Microtask available, send us your CV!

Cristina Sarasua
Institute for Web Science and Technologies (WeST)
University of Koblenz-Landau, Germany
Email: csarasua@uni-koblenz.de

Matthias Thimm
Institute for Web Science and Technologies (WeST)
University of Koblenz-Landau, Germany
Email: thimm@uni-koblenz.de

***Abstract*—Most current microtask crowdsourcing platforms do not exploit the individual expertise of workers, which becomes extremely relevant for knowledge-intensive microtasks in human computation scenarios. In this paper, we discuss work in progress on worker profiling within microtask platforms to increase both the quality of the work and the satisfaction of the users. We analyse the issue of profiling workers and propose the introduction of a crowd worker CV as a comprehensive means to describe a worker's expertise and interests. We discuss several important dimensions that should be included in such a CV and analyse their benefits.**

***Keywords*—Crowd worker, microtask, curriculum vitae, profile, expertise, task assignment**

I. Introduction

The number of people embracing the work published in online labor marketplaces has increased considerably in the last years. Moved by different motivations [1], but with a strong focus on the financial compensation, different people with different cultural and educational backgrounds [2] spend hours contributing to Amazon Mechanical Turk (MTurk) and many other marketplaces. Even though there are some spammers—who instead of providing honest results give deliberately random responses—there is a big workforce who is open to accomplishing microtasks, also in real-time settings. The possibility of having thousands of workers willing to carry out microtasks, even regularly, brings many opportunities to human computation scenarios where humans can aid computers in accomplishing tasks which are not so easily solvable by automatic techniques. In such tasks the expertise of workers can make a difference in both the accuracy of results and the efficiency of the approach. Thus, in order to optimise results it is important to be able to manage such a valuable workforce in a similar manner as for ordinary "office jobs". There, managers have to carefully select their candidates, form suitable teams, and give tasks to the employees who are better qualified and therefore are more likely to solve the tasks with success. Similarly, not all workers in microtask crowdsourcing have the same skills and interests which can, in particular, be observed in the answers they provide to particular microtasks.

The workflow followed in current markeplaces gives workers the freedom to browse the list of available microtasks and select the work they want to do. Requesters can establish basic restrictions and create qualification tests to filter the workers who are allowed to work on their microtasks. But studies as the one of Kittur and colleages [3] suggest that workers would appreciate some help from the platforms in order to reduce the effort they need to invest when looking for appropriate tasks. Now that the set of available microtasks is large and heterogeneous (and can change constantly), there is a need for improving the process of task assignment with worker profiles. Including such a feature in next generation microtask marketplaces would certainly benefit all involved agents, as workers would find suitable microtasks in an easier way and would therefore presumably feel more satisfied, while requesters would get the responses from more reliable workers.

We propose the use of a cross-platform Curriculum Vitae (CV), which similarly to traditional CVs contains information that describes the expertise of the crowd worker and can be used to evaluate whether a worker is a good candidate for a particular task (no matter whether this is used within a process where tasks are assigned to workers or workers to tasks). The crowd worker CV aggregates different types of information which are extracted from different sources (e. g. information extracted from previous accomplished microtasks). In our proposal, the CV is expected to be defined and updated partly by crowd workers, partly by requesters, and partly by (internal or external) automatic techniques which analyse crowd workers.

In the remainder of this short paper we discuss the current situation and recent works on worker profiling in microtask crowdsourcing (Sec. II) and have a look into its possible future (Sec. III).

II. Profiling crowd workers

Profiling and recommendation techniques have proven to be very useful in other areas where personalisation is important, like for example in online shopping environments, tagging systems, and Web search. Initial results indicate that this is a promising research area for crowdsourcing, too.

A. Workers management in current labor marketplaces

One of the characteristics that has been associated to microtask crowdsourcing since the term was coined is anonymity. Usually, workers provide personal data when they register at marketplaces, but this data is not revealed completely to requesters. Requesters who publish microtasks directly at MTurk or via CrowdFlower—two of the most popular microtask platforms at the moment—can obtain the location of the worker, his/her IP, without knowing the identity of the person (i. e. complete name, age, affiliation etc.). Nevertheless, recent research has revealed that MTurk is not as anonymous as many people thought: as Lease et al. explain, the worker identifier, that requesters get from MTurk, can be related to their Amazon account, which—if not configured properly—can provide a lot of personal information [4]. Indeed, this does not contradict MTurk's privacy policies, as the authors of the paper describe. Yet, when workers were informed about this issue during the survey carried out for the paper, workers felt very uncomfortable with the idea of requesters having access to their Amazon wishlist, reviews, pictures, or any other details that reveal their identity.

B. Different ways to obtain information relevant for worker profiling

There are different possible ways to gather relevant information about workers. We classify the different sources of information considering whether the information can be simply accessed or has to be deduced, the type of the source, the type of interaction with the workers, as well as the granularity of the task performance analysis.

1) Explicit information: information that can be directly accessed because it was entered explicitly.
   a) Inside the marketplace: information that is provided within the marketplace where the microtask has been published.
      i) Registration information: information that was entered by the worker on registration, such as demographic information and gender.
      ii) Qualification tests or requested feedback: information that is obtained after asking the worker to disclose his/her expertise, e. g. whether the worker can translate technical manuscripts or whether he/she has some specific domain knowledge.
   b) Outside the marketplace: information that is available in a third party source. For example, the "likes" published on Facebook, or the reviews given at Amazon.
2) Implicit information: information that has not been communicated directly, but can be deduced.
   a) Among different types of microtasks: information that can be learned comparing the behaviour in different types of microtasks. For example, the performance obtained in annotation microtasks, in data interlinking microtasks and text translation microtasks.
   b) Within the same type of microtask: information that can be deduced comparing the behaviour in different particular units or instances of microtask of the same kind of microtask. For example, the accuracy obtained in interlinking data that refers to the US, and data that refers to Europe.

We continue with discussing some case studies that take the above mentioned sources into account for personalising microtask crowdsourcing.

C. Our own experience: crowdsourcing data interlinking

In [5] we investigated whether data interlinking on the Web of Data can be effectively outsourced to human crowd workers and whether considering worker profile information increases the accuracy of results. Roughly, the data interlinking problem is about recognising that two *resources* of different data sets refer e. g. to the same real-world entity. For example, imagine some data set A containing information about British authors and their works, and another data set B about policy modelling and voting rules. In A there may be some entity "Lewis Carroll" (the author of e. g. *Alice in Wonderland*) and in B there may be some entity "Charles Lutwidge Dodgson" (the inventor of the Dodgson voting rule). Surprisingly, both entities refer to the same real-world entity but this can hardly be recognised by automatic systems nor by humans not familiar with neither of these two names. Most people would presumably say that these two names do not refer to the same real-world entity which might give wrong results when relying on simple aggregation strategies for crowdsourcing such as majority voting. In CrowdLINK—our approach for crowdsourced data interlinking—the profile of a worker is learned by observing his/her performance in previous tasks. For that, the individual interlinking tasks are categorised depending on e. g. their topic (events, persons, etc.) or their location (US, Europe, etc.), and subsequent tasks are assigned to a worker depending on his/her performance of other tasks from the same or similar categories. Although the presented formal model for the profiling and recommendation components are rather simple (see [5] for the technical details) the first results were quite encouraging: taking profile information into account increased the recall—i. e. the number of correctly identified links—by up to 35 %.

D. Other initiatives which analyse worker profiles

The approach of Bozzon et al. [6] also considers the issue of finding experts to work on specific tasks. In contrast to [5] the approach of [6] does not observe the task performance of the workers but takes information gathered from social networks into account. Using text matching techniques they search through the friend/follower network of a worker to discover content items such as tweets or posts related to some given information need (as e. g. needed for solving a crowdsourcing task). By employing standard information retrieval techniques they rank workers by their expertise to satisfy the given information need. Bozzon et al. tested their approach by gathering the profile information from Facebook, Twitter, and LinkedIn. In general, they discovered that taking more information into account—in particular, information about friends from the user's networks—is beneficial to assess the expertise of a user. Surprisingly, Twitter seems to be the best resource for determining whether a given user has the needed expertise while LinkedIn is less helpful. Similarly, Difallah et al. [7] also use information extracted from social networks to implement a push strategy for crowdsourcing platforms, i. e. a recommendation strategy for crowdsourcing tasks. In particular, they use Facebook to gather both information from the users and also as a platform for the crowdsourcing prototype (via a Facebook App). In their experimental study they show that (some of) their approaches outperform the standard pull methodology of crowdsourcing platforms by a relative improvement of up to 29 % in accuracy.

In general, using social networks to gather relevant information about users seems useful as has also been shown e. g. by studies on predicting personality traits from Facebook activity [8]. The work [9] also uses personality traits predicted from task completion to categorise workers into five different classes and evaluate their performance with respect to those classes. In [9] the relationship between personality traits and the successful completion of tasks from a relevance labelling scenario have been investigated. One particularly interesting result of [9] is that successful completion of such a labelling task is strongly correlated to the Openness trait (the trait describing among others a person's curiosity and creativity).

The work [10] discovered that workers revealing demographic information outperform anonymous workers in a simple word-counting task scenario, in particular when cooperation between different workers is needed. In [11] Satzger et

Approach	1.a.i	1.a.ii	1.b	2.a	2.b
Sarasua et al. [5]					X
Bozzon et al. [6]			X		
Difallah et al. [7]			X		
Kazai et al. [9]		X			X
Huang and Fu [10]	X				
Satzger et al. [11]		X		X	
Khazankin et al. [12]		X			
Amati et al. [13]		X		X	

TABLE I. COMPARING APPROACHES

al. introduce auction-based task assignment for crowdsourcing scenario where workers bid for tasks they like to solve. Besides this they also include a limited form of learning a worker's profile by comparing his/her performance against users with known profile information. Khazankin et al. [12] investigate specifically the issue of scheduling in crowdsourcing when taking expertise into account, while Ambati et al. [13] describe work in progress on learning user profiles from the interactions on microtasks.

In Table I we compare these recent approaches with respect to the way they obtain information about the workers, as it has been discussed in Sec. II-B.

III. LOOKING INTO THE FUTURE

As shown in the previous section, there are several types of information that can provide a useful insight into the expertise of workers. Far from being alternatives, these pieces of information are complementary because they describe different aspects of workers. Then, why not aggregate all the information and build a Curriculum Vitae (CV) to report the worker expertise as we do in a traditional job setting?

A. Defining a CV for crowd workers

CrowdFlower allows requesters to publish microtasks in more than 30 marketplaces at the same time. This is certainly a positive feature, as publishing in several places at the same time can make the completion time shorter. However, when the results are retrieved from different marketplaces there is a loss of information: if a worker has submitted a response to two instances of the the same kind of microtask (i. e. two microtasks generated by CrowdLINK), CrowdFlower is not able to identify that the same person was carrying out both microtasks. If this information was available, the analysis of the worker performance would be more accurate and requesters could better control that a microtask is done only once by the same worker. This motivates the requirement of the CV to be cross-platform, representing one unique worker in the real world. Each worker could have accounts in different marketplaces and only one CV.

Table II shows the kind of information that a crowd worker CV should contain. Taking into account the information tackled by some of the existing initiatives, and extending it with our own suggestions, we have divided the CV into six sections: personal data and demographic information, interests, skills, professional experience, feedback and requester/marketplace evaluation. For each section, the table shows the information we would like to keep in the CV, who (and how) should provide this information and whether it should be optional (Opt.) or mandatory to collect it. Note that we deliberately exclude specific data about the identity of the worker (e. g. complete name), but it would be possible to include this data,

Feature	Definition	Who/How	Opt.
Personal data and demographics			
Marketplace identifiers	The list of the different identifiers that the worker has for different marketplaces	Worker, manually	yes
Location	Name of country and city of the worker	Worker, manually	no
Interests			
Extracted topics	The list of topics the worker is interested in	Automatically extracted from social networks	yes
Stated topics and types of microtasks	The list of topics and types of microtasks that the worker states as of his/her interest	Worker, manually	yes
Skills			
Languages	The natural languages known by the worker	Worker, manually	no
Qualification tests	List of results of the qualification tests that the worker answered	Marketplace, automatically	no
Professional experience			
Accomplished tasks	Number and type of accomplished microtasks (total / AVG) in general and per type of microtask, together with the performance of the worker	Marketplaces, automatically	no
Work frequency	How frequently a worker is available in marketplaces	Marketplaces, automatically	no
Time	AVG time spent by a worker in a particular type of microtask	Marketplace, automatically	no
Declared confidence	The AVG confidence a worker says (s)he has in a particular kind of microtask	Worker, manually	yes
Earned rewards	The amount of money that a worker received per type of microtask and marketplace	Marketplaces, automatically	no
Bonus	The amount of money bonus that a worker received per type of microtask and marketplace	Marketplaces, automatically	no
Feedback			
Comments	List of comments given by the worker	Worker, manually	yes
Satisfaction with marketplace	The score that a worker gives to a marketplace after working at it	Worker, manually	yes
Requester/marketplace evaluation			
Flags	Number of negative acquired flags	Requester, manually	no
Reputation / global accuracy	The score of the worker's accuracy in each marketplace	Marketplace, automatically	no

TABLE II. RELEVANT INFORMATION FOR THE CROWD WORKER CV

too. Because of privacy issues we define features like the IDs of the worker in each marketplace as optional, as there might be workers who are concerned with the fact of being tracked across different platforms. The more information the CV contains, the more accurate the profile will be and the better assessments will be done in task assignment.

As it happens in traditional CVs, the specified information should be constantly updated. The CV will be defined in a centralised location, so that when a new worker starts using a marketplace for the first time a CV is generated for him/her, containing the information manually provided by the worker at registration time. It can be implemented as an application, so that the CV becomes active and updated when the worker logs into a marketplace. As soon as the worker provides more information and carries out microtasks, the CV will be updated either by the marketplaces, the requesters or other automatic means for analysing external sources (e. g. social networking sites). A way to control possible CV spammers should be defined.

In order to provide machine-processable CV data that marketplaces can exchange, we envisage the definition of an (extendable) RDF[1] vocabulary (called CrowdWorkerCV) to describe the information represented in Table II. The vocabulary could take into account the ResumeRDF ontology, and other microformats like hResume and vCard[2], which are able to describe professional user data. However, the CrowdWorkerCV vocabulary would mainly define things which are specific to microtask crowdsourcing (e. g. the number of microtasks accomplished), as the reader can observe from the content of the table.

B. Discussion

When we analyse the influence that a crowd worker CV can have on both workers and requesters, we identify several aspects that can be seen as advantages and disadvantages.

- *Workers*: a positive aspect for workers is that a detailed CV enables recommendation of microtasks, which saves a lot of time to crowd workers. Along the same lines, with appropriate microtasks workers can hopefully feel more comfortable and satisfied. The less positive aspect is that some workers might be concerned with privacy issues and would not agree to provide data such as their identifiers in different marketplaces, even if this would not necessarily reveal their identity. However, as mentioned before, we discourage the mandatory definition of a complete CV. Workers would have the freedom to decide the "sensitive" information they like to provide.
- *Requesters*: if requesters can have access to the aforementioned information, they will be able to identify which workers are using several platforms and thus the performance will be analysed across the different marketplaces. The possibility of addressing the most reliable worker will improve accuracy and efficiency of results, and the awareness of being analysed can make workers work more strictly which will also improve the quality of results. The fact of making the CV globally accessible can be a disadvantage for some requesters, or even for the marketplace owners, if they feel their marketplace should be considered differently from others due to own business rules. Therefore, the CV should not interfere with separate analyses, but should still reflect the shared view across microtask platforms which aims to benefit all agents. Our CV can be benefitial for marketplaces, as they can use the information about the workers to compete against other marketplaces (e. g. they can lure more requesters/microtasks claiming they have the workers with the best performance in a particular type of task). This would mean a higher income for them.

So, overall a crowd worker CV is likely to have a positive impact. It can make the analysis of workers more precise and richer, something that will benefit both workers and requesters.

IV. Conclusions

It is not clear whether microtask crowdsourcing will evolve and encourage more transparency between workers and requesters, as well as among workers. What we expect (by observing other kinds of online systems where recommendation was applied) is that revealing the identity will ultimately depend on workers who will consciously allow or disallow that others know who they are. This could lead to a situation where some workers with a more complete profile and other workers with a less descriptive profile both reside in the same marketplace. We foresee that the challenge in such a case will be to distribute the total set of workers in a fair way (without mandatorily marginalising workers whose identity is unknown), still optimising the match between workers and microtasks so as to let requesters have the best results possible.

Even if anonymity is preserved, identifying worker profiles can be beneficial for all microtask crowdsourcing players. Specially for microtasks where the lack of background knowledge can reduce the probability of guessing the correct answer considerably, we need to ensure that the appropriate workers are addressed. Using a cross-platform CV, defined as RDF data, we can aggregate all the relevant information to describe a worker's expertise, so as to facilitate the evaluation of workers among different marketplaces. Future work will include the formal definition of the CrowdWorkerCV vocabulary, and the study of the acceptance of such a CV (both by workers and requesters).

[1] http://www.w3.org/TR/rdf-syntax-grammar/

[2] ResumeRDF http://rdfs.org/resume-rdf/, hResume http://microformats.org/wiki/hresume, vCard http://microformats.org/wiki/vcard

References

[1] N. Kaufmann and T. Schulze, "Worker motivation in crowdsourcing and human computation," in *Proceedings of the AAAI workshop on human computation (HCOMP)*, 2011.

[2] P. G. Ipeirotis, "Demographics of mechanical turk," in *NYU Working Paper No. CEDER-10-01.*, 2010.

[3] A. Kittur, J. V. Nickerson, M. S. Bernstein, E. M. Gerber, S. Aaron, J. Zimmerman, M. Lease, and J. J. Horton, "The future of crowd work," in *16th ACM Conference on Computer Supported Coooperative Work (CSCW 2013)*, 2013.

[4] M. Lease, J. Hullman, J. P. Bigham, M. S. Bernstein, J. Kim, W. Lasecki, S. Bakhshi, T. Mitra, and R. C. Miller, "Mechanical turk is not anonymous," 2013, sSRN.

[5] C. Sarasua, M. Thimm, S. Staab, and E. Simperl, "Data interlinking using crowd workers with background knowledge," 2013, submitted.

[6] A. Bozzon, M. Brambilla, S. Ceri, M. Silvestri, and G. Vesci, "Choosing the right crowd: expert finding in social networks," in *Proceedings of the 16th International Conference on Extending Database Technology (EDBT2013)*, 2013.

[7] D. E. Difallah, G. Demartini, and P. Cudré-Mauroux, "Pick-a-crowd: tell me what you like, and I'll tell you what to do," in *Proceedings of the 22nd international conference on World Wide Web (WWW2013)*, 2013.

[8] Y. Bachrach, M. Kosinski, T. Graepel, P. Kohli, and D. Stillwell, "Personality and patterns of facebook usage," in *Proceedings of the 3rd Annual ACM Web Science Conference*, 2012.

[9] G. Kazai, J. Kamps, and N. Milic-Frayling, "Worker types and personality traits in crowdsourcing relevance labels," in *Proceedings of the 20th ACM international conference on Information and knowledge management*, 2011.

[10] S. Huang and W. Fu, "Dont hide in the crowd! increasing social transparency between peer workers improves crowdsourcing outcomes," in *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, 2013.

[11] B. Satzger, H. Psaier, D. Schall, and S. Dustdar, "Auction-based crowdsourcing supporting skill management," *Information Systems, Elsevier*, 2012.

[12] R. Khazankin, H. Psaier, D. Schall, and S. Dustdar, "QoS-based task scheduling in crowdsourcing environments," in *Proceedings of the 9th international conference on Service-Oriented Computing*, 2011.

[13] V. Ambati, S. Vogel, and J. G. Carbonell, "Towards task recommendation in micro-task markets." in *Human Computation*, 2011.